THE VERTICAL STRUCTURE OF THE OUTER MILKY WAY HI DISK

E.S. Levine[1], Leo Blitz[1], and Carl Heiles[1]
Department of Astronomy, University of California at Berkeley, Mail Code 3411, Berkeley, CA 94720 USA
Accepted by ApJ — 29 Jan 2006

ABSTRACT

We examine the outer Galactic HI disk for deviations from the $b = 0°$ plane by constructing maps of disk surface density, mean height, and thickness. We find that the Galactic warp is well described by a vertical offset plus two Fourier modes of frequency 1 and 2, all of which grow with Galactocentric radius. Adding the $m = 2$ mode accounts for the large asymmetry between the northern and southern warps. We use a Morlet wavelet transform to investigate the spatial and frequency localization of higher frequency modes; these modes are often referred to as "scalloping." We find that the $m = 10$ and 15 scalloping modes are well above the noise, but localized; this suggests that the scalloping does not pervade the whole disk, but only local regions.

Subject headings: Galaxy: disk — Galaxy: structure — Galaxy: kinematics and dynamics — ISM: structure — radio lines: general

1. INTRODUCTION

Although the topography of the gas disk of the Milky Way has previously been mapped (Westerhout 1957; Henderson et al. 1982; Burton & Te Lintel Hekkert 1986, and many others), its complete Fourier structure has never been quantitatively described, though Binney & Merrifield (1998) have approximated its three lowest frequency terms. Which spatial oscillation frequencies are most important? Is the scalloping a local or global effect? This paper constitutes an in depth analysis of the shape of the outer HI disk, as well as the first quantitative analysis of the scalloping.

The large-scale warp in the gas disk of the Milky Way has been known since 1957 (Burke 1957; Kerr 1957; Westerhout 1957; Kerr et al. 1957). The warp has a large amplitude, rising to a height greater than 4 kpc at a Galactocentric radius of 25 kpc in the northern data. It is also asymmetric; in the south the gas falls about 1 kpc below the plane before rising back to it. By observing other galaxies, it is possible to develop a general understanding of how warps behave. Bosma (1991) found that at minimum half of all galaxies are warped, and that galaxies with smaller dark matter halo core radii are less likely to be warped. Briggs (1990) claimed that a warp's line of nodes starts out straight, and at a transition begins to advance in the direction of rotation, with some exceptions. Another survey has shown that warps are common in galaxies with HI disks that are extended compared to their optical components, and are often asymmetric (García-Ruiz et al. 2002b). In the Milky Way, many components other than HI also participate in the warp. A partial list includes dust (Freudenreich et al. 1994), CO (Wouterloot et al. 1990), solar neighborhood stars (Dehnen 1998), and IRAS point sources (Djorgovski & Sosin 1989).

Many efforts have been directed toward understanding the warp on a theoretical basis. Bending modes have long been suspected as the mechanism creating and maintaining the warp. Early work studying the evolution of bending mode oscillations was hampered by a lack

of knowledge regarding galactic halos, but showed that the shape of the density falloff near the edge of the disk plays an important role in the stability of bending modes (Hunter & Toomre 1969; Toomre 1983). The distribution of matter in the halo controls the ability of the disk to sustain long-lived bending wave warps, so studying the properties of the warp will tell us about the shape of the halo (Binney 1978; Sparke 1984). Sparke & Casertano (1988) argued that bending mode oscillations are plausible when the halo and the self-gravity of the disk are taken into account, but Binney et al. (1998) showed that in this situation the warp will wind up within a few dynamical times. Previous work has largely been concerned with the existence and behavior of $m = 1$ warps only, though Sparke (1995) demonstrated the stability of $m = 0$ modes in an axisymmetric halo.

Several other mechanisms have been suggested as possibilities for creating and maintaining a warp (Kuijken & García-Ruiz 2001). Gravitational interaction with satellites such as the Magellanic Clouds is a promising candidate (Burke 1957; Weinberg 1998; Weinberg & Blitz 2005), but there is a longstanding debate of whether tidal effects are strong enough to produce the observed effect (Kerr 1957). Indeed, even in galaxies with companions similar to the LMC, tidal amplification may not be strong enough to account for the size of the warp (García-Ruiz et al. 2002a). Accretion of matter onto the halo is another plausible cause (Jiang & Binney 1999), or matter can accrete directly onto onto the disk and torque the gas orbits (López-Corredoira et al. 2002; Sánchez-Salcedo 2004). Intergalactic magnetic fields can act on a slightly ionized gas disk to produce a warp (Battaner et al. 1990), or the intergalactic medium can excite a warp through a wind (Kahn & Woltjer 1959).

Throughout this paper, we will refer to shorter wavelength (5–25 kpc) oscillations in azimuth as scalloping. Gum et al. (1960) first noticed a "waviness" in the gas layer of the inner Galaxy; observational evidence for scalloping turns out to be prominent in the outer Galaxy (Henderson et al. 1982; Kulkarni et al. 1982). Fleck (1983) investigated the possibility that the Milky Way scalloping results from the Kelvin-Helmholz insta-

bility, and Sparke (1995) suggested that the scalloping should only be present in the outer parts of Galactic disks. Spicker & Feitzinger (1986) find azimuthal corrugations in HI and other components along the spiral arms after removing the warp. In an N-Body simulation, Edelsohn & Elmegreen (1997) found evidence for spiral corrugations due to gravitational interaction with a satellite galaxy. In this paper we will only investigate oscillation in Galactocentric azimuth, and not in radius.

In §2 we describe the method of transforming the Heliocentric data cube into Galactocentric coordinates, and present maps of the surface density, average height off the $b = 0°$ plane, and vertical thickness of the outer Galaxy. While these maps are constructed using new data, they are not significantly different from previous work. In §3 we perform global and local analyses on these maps to better understand the warp and the scalloping.

2. METHOD

2.1. Data Processing

We use the 21 cm Leiden/Argentine/Bonn (LAB) data (Kalberla et al. 2005; Hartmann & Burton 1997; Bajaja et al. 2005; Arnal et al. 2000) to conduct a quantitative study of the warp and scalloping. The LAB survey is a combination of the LDS data set (Hartmann & Burton 1997) with Southern sky observations from the IAR (Arnal et al. 2000); however, much attention has been paid toward ensuring a uniform data set. The data are corrected for stray radiation. The combined survey maps the entire sky within $-450 \le v_r \le 400$ km s^{-1} with a resolution of 1.3 km s^{-1}; this velocity range includes all of the gas in the Galaxy in circular rotation. We used the Hanning smoothed data, which have a velocity resolution of 1.9 km s^{-1}, and used only data with $|b| \le 30°$. The signal of the warp in the north can be weakly traced beyond this elevation limit (Burton 1985), but the vast majority of the warp signal is included within our b range.

The LAB survey contains a large number of angularly small emission features, especially at high latitudes, which are not contiguous with the gas in the disk, and may not even be in circular rotation. Features like these are particularly troublesome at larger Galactic latitudes since they can contain enough gas to contaminate our calculations, especially at large Galactocentric radii. These objects are removed with a median filter so we can focus on the gas in the disk. Points with brightness temperature $T_b \le 0$ are temporarily filled with the value 0.01 K for this filtering only. We then cycle through the LAB data cube and calculate the median of each point and its 12 nearest neighbors in ℓ and b at the same line-of-sight velocity v_r (a two dimensional diamond shaped filter); call this median $T_m(\ell, b, v_r)$. Any point with $T_b > 10\,T_m$ is replaced with $10\,T_m$. Two examples of objects caught by this filter are M31 and NGC 6822.

The survey grid in ℓ, b, and v_r is not convenient for analyzing Galactic properties. Ideally, our measurements would be equally spaced in the Galactocentric cylindrical coordinates R, ϕ, and z; we must interpolate a Galactocentric grid from the LSR-centered data. The Galactic azimuth ϕ is defined such that it converges with galactic longitude ℓ at large R. To convert from observed ℓ, b, and v_r to R, ϕ, and z we use the following mapping

functions:

$$\ell = \sin^{-1}\left[\frac{R}{r'}\sin\phi\right]$$

$$b = \tan^{-1}\frac{z}{r'}$$

$$v_r = \sin\ell\cos b\left[\frac{R_0}{R}\Theta(R) - \Theta_0\right]$$

$$+ v_\Pi(R)\cos\phi\left(1 - \frac{R_0^2}{R^2}\sin^2\ell\right)^{1/2}\cos b. \quad (1)$$

Here, $\Theta(R)$ is the Galactic rotation curve, which we assume to be 220 km s^{-1} everywhere (Brand & Blitz 1993). Θ_0 and R_0 are 220 km s^{-1} and 8.5 kpc, respectively. **r** is the vector connecting the Sun's location to the point under investigation; **r'** is the projection of this vector onto the plane of the disk (with magnitude r').

Equation 1 is the transformation that results from an assumption of elliptical gas orbits with major axis along the $\phi = 90°, 270°$ line. To derive this, we assume the gas moves on the orbit:

$$v_\phi = 220 \text{ km s}^{-1}$$

$$v_R = v_\Pi(R)\cos\phi. \quad (2)$$

v_R is the magnitude of the velocity in the Galactocentric radial direction. v_Π is a parameterization of the ellipticity of the orbit, which is free to vary with Galactocentric radius; we discuss our method of calculating $v_\Pi(R)$ in the Appendix. At all points, $v_R/v_\phi < 0.1$. These equations are simply the epicyclic approximation for an orbit with epicyclic frequency 1 and the angle of ellipse orientation fixed. Although the gas orbits in the Galaxy are not likely to correspond to the fixed ellipse orientation we describe, this configuration minimizes the correction to v_r for gas far from $\ell = 0°$ or 180° (see Voskes (1999) for more detail on elliptical gas orbits). Without a correction for radial motion of the gas, there is a large asymmetry between the surface densities at Galactic longitudes on either side of $\ell = 0°$ and $\ell = 180°$ (Henderson et al. 1982). This must be taken into account, or features in these two regions such as the surface density will appear discontinuous and distorted. Assuming an outward velocity for the Local Standard of Rest (LSR) will correct this discontinuity to some degree (Kerr 1962; Kerr & Westerhout 1965). However, it seems that the best fit for the motion relative to the LSR changes with radius, implying that the effect is global, rather than local (Blitz & Spergel 1991). To reduce the magnitude of the discontinuity using gas orbits, one can use a Galactocentric radial velocity roughly of the form $\cos\phi$ (or $\cos\ell$) (Kuijken & Tremaine 1994).

We exclude all points that lie within $345° \le \ell \le 15°$ or $165° \le \ell \le 195°$. Points in these two wedges have velocities along the line of sight that are too small with respect to their random velocities to establish reliable distances. All points in this region are set to $T_b = 0$.

Using (1) we construct a Galactocentric grid $T_b(R, \phi, z)$ by trilinear-interpolating from the grid $T_b(\ell, b, v_r)$. We do this by calculating the coordinates of an (R, ϕ, z) point in (ℓ, b, v_r) space, and interpolating from $T_b(\ell, b, v_r)$. The resolution of the Galactocentric grid is set by the spacing of the LSR centered grid, but in this paper we are not interested in small-scale disk structure. A grid of 100

points in 10 kpc $\leq R \leq$ 30 kpc, 350 points in $-\pi \leq \phi \leq \pi$, and 141 points in -20 kpc$\leq z \leq$ 20 kpc gives us sufficient resolution to answer the questions we are interested in.

Undersampled grids fail to utilize all of the information in the data; our choice of grid spacing is both is an undersampling and an oversampling of the information in $T_b(\ell, b, v_r)$ depending on the position in the disk. Consider two points near $R = 10$ kpc and $\ell = 15°$, where we have poor Galactocentric resolution in ϕ. The LAB survey has $\Delta\ell = 0.5°$; at this location this corresponds to $\Delta\phi \approx 0.9°$. The spacing in our Galactocentric grid is larger than $1°$, thus we are undersampled everywhere in the ϕ dimension. In §3.1.2 we will perform an azimuthal frequency analysis of each ring; none of the frequencies we examine approach the Nyquist frequency of the data. Near the midplane at $R = 30$ kpc and $\ell = 15°$, lines of constant b are separated by around 300 pc, and near the top of the grid they are separated by around 400 pc; there is some oversampling in the z dimension by no more than a factor of 2. The most severe case of oversampling is in the R dimension, where along $\ell = 15°$ we have only 14 velocity resolution elements in our R range for the 100 grid points. On the other hand, at $\ell = 90°$, there are more than 60 resolution elements in our range. We have chosen a grid spacing in R that oversamples to a varying degree depending on location in the disk.

From $T_b(R, \phi, z)$, we can recover $\rho(R, \phi, z)$ using the method outlined in Kerr (1968). We assume T_s, the spin temperature, is 155 K everywhere. This is slightly higher than the maximum brightness temperature found in the LAB survey in the region we are concerned with. Choosing a larger number to force the optically thin limit ($T_b \ll T_s$) makes a difference only in a small number of areas in the inner radii of our grid. The vast majority of the points are optically thin with any reasonable T_s, and are not affected by this choice. The transformation to a density grid depends on $|dv_r/dr|$. Using elliptical orbits makes calculating this quantity slightly more difficult than with a flat rotation curve. Although the full derivative can be written analytically, we just calculate it numerically. Points with $T_b < 0$ are set to $\rho = 0$.

2.2. Surface Density and Mean Height maps

The grid $\rho(R, \phi, z)$ contains information about the density of HI in the Galaxy, minus whatever has been removed by the median filter and the excluded regions. Previous studies have proceeded by calculating a mean height $\bar{z}(R, \phi)$ for the gas. However, the Galaxy is a complicated place that contains a variety of HI structures in addition to the disk. In particular, there are many extended clouds located near the disk as well as spurs that split off from the disk. None of these will have been removed by the median filter, which acts only on comparatively small areas of the sky. Voskes (1999) developed an alternative method to remove some of these features; he masked out a map of high velocity cloud complexes. Since we are only interested in the shape of the disk itself, these additional components must be filtered out before our calculation of the mean height. We perform a dispersion filter that operates as follows.

1. Calculate the total mass surface density

$$M(R, \phi) = \sum_{i=1}^{N=141} \rho(R, \phi, z_i)\Delta z \quad (3)$$

where Δz is the z bin size.

2. Calculate

$$\bar{z}(R, \phi) = \frac{\sum_{i=1}^{N=141} z_i \rho(R, \phi, z_i)\Delta z}{M(R, \phi)}. \quad (4)$$

3. Calculate the second moment

$$d^2(R, \phi) = \frac{\sum_{i=1}^{N=141} [z_i - \bar{z}(R, \phi)]^2 \rho(R, \phi, z_i)\Delta z}{M(R, \phi)}. \quad (5)$$

4. Run through each point in (R, ϕ, z) space. For any point that does not lie within $2d$ of $\bar{z}$, set $\rho = 0$. Call the grid that results from this dispersion filter $\rho_d(R, \phi, z)$.

For step 4, we experimented with several different cutoffs (in the range $1 - 3d$); our results do not depend strongly on which cutoff we choose.

We construct the Galactic disk surface density and mean height maps from $\rho_d(R, \phi, z)$. For example, we can sum the density over the z dimension to construct the surface density associated with the disk (Figure 1):

$$\Sigma(R, \phi) = \sum_{i=b}^{t} \rho_d(R, \phi, z_i)\Delta z \quad (6)$$

The indices t and b represent the top and bottom z_i that have not been zeroed out by the dispersion filter. The resulting figure clearly demonstrates the falloff of the disk surface density with radius. The contour lines are nicely continuous across the $\ell = 0°$ and $\ell = 180°$ lines because of our use of elliptical orbits; see the Appendix for a version of this figure without correcting for these orbits. The jagged nature of the contours with $R \lesssim 18$ kpc is due in part to spiral arms (Levine et al. 2006). There is a region near $R \approx 27$ kpc and $\phi \approx 90°$ with a smaller surface density than other regions at the same radius; this region has somewhat unusual features in all of our maps. Also notice the excluded regions near the Sun-Galactic center line; these gaps will appear in all of our plots.

We will be looking closely at the mass weighted mean height of the gas disk (Figure 2),

$$h(R, \phi) = \frac{1}{\Sigma(R, \phi)} \sum_{i=b}^{t} z_i \rho_d(R, \phi, z_i)\Delta z. \quad (7)$$

This height is calculated with respect to the Galactic midplane defined by $b = 0°$. The Galactic warp is the most immediately evident feature in this map; the gas in the northern hemisphere peaks at $h \approx 5$ kpc, while the southern gas descends only to $h \approx -1.5$ kpc, consistent with previous maps of the Galaxy. At least three vertical oscillations of magnitude ≈ 1 kpc can be seen in the south from $\phi \approx -120°$ to $\phi \approx -20°$; these have previously been called the "scalloping". The approximate extent of the scalloping is marked with an arc connecting two "S" labels. The region of low surface density noted

in the discussion of Figure 1 ($R = 27$ kpc, $\phi = 90°$) has a height that seems anomalous when compared to surrounding gas (it is marked with an "X"). Several features are elongated along lines of constant ℓ indicating some level of contamination by turbulent velocities and/or local gas. In the past these have been dubbed "fingers of God" because they all point back to the Sun.

In §3.1.1 we will need a measure of the uncertainty in the mean height of each point in the map. We define this error using the sum of squared residuals, as is usual for least squares fits:

$$e^2(R, \phi) = \frac{\sum_{i=b}^{N=t}(z_i - h(R, \phi))^2 \rho_d(R, \phi, z_i)\Delta z}{(t - b + 1)\Sigma(R, \phi)}, \quad (8)$$

where $t - b + 1$ is the number of points in the calculation of the mean and surface density. This is only an approximation because it does not account for any of the uncertainty introduced in $T_b(\ell, b, v_r)$, $v_\Pi(R)$, the interpolation to $\rho(R, \phi, z)$, or the dispersion filter. Typically, $e(R)/R \approx 0.01$.

Because of effects like turbulence, anomalous velocities, and spiral arms, the features in the mean height map may not correspond to the actual shape of the Galaxy. The severity of this effect can be seen by looking at a contour map of dr/dv_r (see Figure 3). While similar to previously published plots (Burton & Te Lintel Hekkert 1986), this figure also includes the correction for elliptical gas orbits. Given a turbulent velocity of magnitude v_t, features with coherent scales less than $v_t dr/dv_r$ could potentially be false signals. Furthermore, even real features will be blurred out or even incorrectly positioned over the same length scale. A typical turbulent velocity is around 8 km s^{-1}. Small regions can differ from the flat rotation curve by 20-30 km s^{-1} (Brand & Blitz 1993). We do not expect that these distortions will significantly effect our Fourier analysis. The steep contours near $\ell = 0°$ and $\ell = 180°$ demonstrate the need for the excluded regions, because small irregularities in velocity there result in large changes in distance.

2.3. Thickness map

We also construct a measure of the thickness of the disk. In contrast to the previous section, we keep all of the complicated features of the HI in our calculation of the thickness. We do this because these features do contain information regarding pressure and gravitational force, although they are a nuisance when calculating the mean height. In particular, our method for finding the thickness of the disk relies heavily on the tails of the vertical density distribution. Thus, in this section, we will work with ρ instead of ρ_d.

Following Henderson et al. (1982), we define the first and third quartile points as z_{j1} and z_{j3} as the smallest and largest indices, respectively, that satisfy:

$$\sum_{i=0}^{j1} \rho(R, \phi, z_i)\Delta z \geq M(R, \phi)/4$$

$$\sum_{i=j3}^{141} \rho(R, \phi, z_i)\Delta z \geq M(R, \phi)/4. \quad (9)$$

We then define the half thickness (a factor of 2 smaller than Henderson et al. (1982) to ease comparison with recent work): $T_h(R, \phi) = (z_{j3} - z_{j1})/2$. Note that this implies that T_h is quantized by $\Delta z/2$ (about 140 pc). This

can lead to inaccuracies in T_h in places where the thickness is small, i.e. $R \approx 10$ kpc. $T_h(R, \phi)$ is shown in Figure 4. Some points near the sun actually have calculated half-thicknesses of zero; this is due to the comparatively small thickness of the disk in that region combined with our limited LAB survey range $|b| \leq 30°$ and poor grid resolution in z.

The flaring of the disk with radius is immediately evident, as T_h increases in magnitude by a factor of about 8 between $R \approx 10$ and 30 kpc; this flaring was first seen in Lozinskaya & Kardashev (1963). Asymmetry between the northern and southern halves of the disk is also prominent; the southern half of the Galaxy has a lower average thickness than the northern. This asymmetry was also evident in Henderson et al. (1982) and Burton & Te Lintel Hekkert (1986). The low surface density region around $R \approx 27$ kpc has a very large thickness; the gas in the region has clearly been disturbed and dispersed by some mechanism. There are regions of increased thickness close to $\ell = 15°$ and $345°$ near $R \approx 10$ kpc probably caused by local gas.

Figure 5 plots the surface density and half-thickness of the HI layer averaged over ϕ as a function of R. This figure confirms our visual impressions from the surface density and dispersion maps by showing the falloff in the surface density and the flaring of the thickness with radius. The contamination by local gas described in the previous paragraph as well as the poor z resolution problem mentioned at the beginning of the section account for the rise in the average of T_h with decreasing R near $R \approx 10$ kpc; there is no evidence the thickness of the disk actually behaves this way. A gradient-expansion least-squares routine (Marquardt 1963) gives the best fit for the surface density beyond 14 kpc (where the exponential falloff begins) as:

$$\Sigma(R) = 4.5 \times \exp[-(R - 14\,\text{kpc})/4.3\,\text{kpc}]\,\text{M}_\odot\,\text{pc}^{-2}. \quad (10)$$

3. ANALYSIS

3.1. Global Structure

Using the maps we have constructed, we conduct a quantitative investigation of the disk vertical structure. Since we are studying a disk, we use a method that is independent of rotation in ϕ and treats $\phi = 0$ and $\phi = 2\pi$ as the same point. Furthermore, the data are unevenly sampled due to the excluded regions within $15°$ of the Sun-Galactic center line. We complete the analysis without extrapolating $h(R, \phi)$ in these regions, to avoid introducing any artifacts into the signal. Throughout this paper, we refer to different frequency oscillations in the disk. These frequencies will always be labeled by the number of oscillations they will complete in a full 2π; thus the $m = 1$ mode has a one maximum and one minimum in the disk.

3.1.1. Galactic Warp

The Galactic warp is the most prominent feature in Figure 2. A Lomb periodogram analysis of each radius ring (see §3.1.2) reveals that the power in each of the $0, 1,$ and 2 modes is consistently larger than that in any other mode for $R \gtrsim 20$ kpc. At some radii $m = 3$ is the next strongest mode, and at others it is $m = 4$. Accordingly, we characterize the warp by an offset in the z direction,
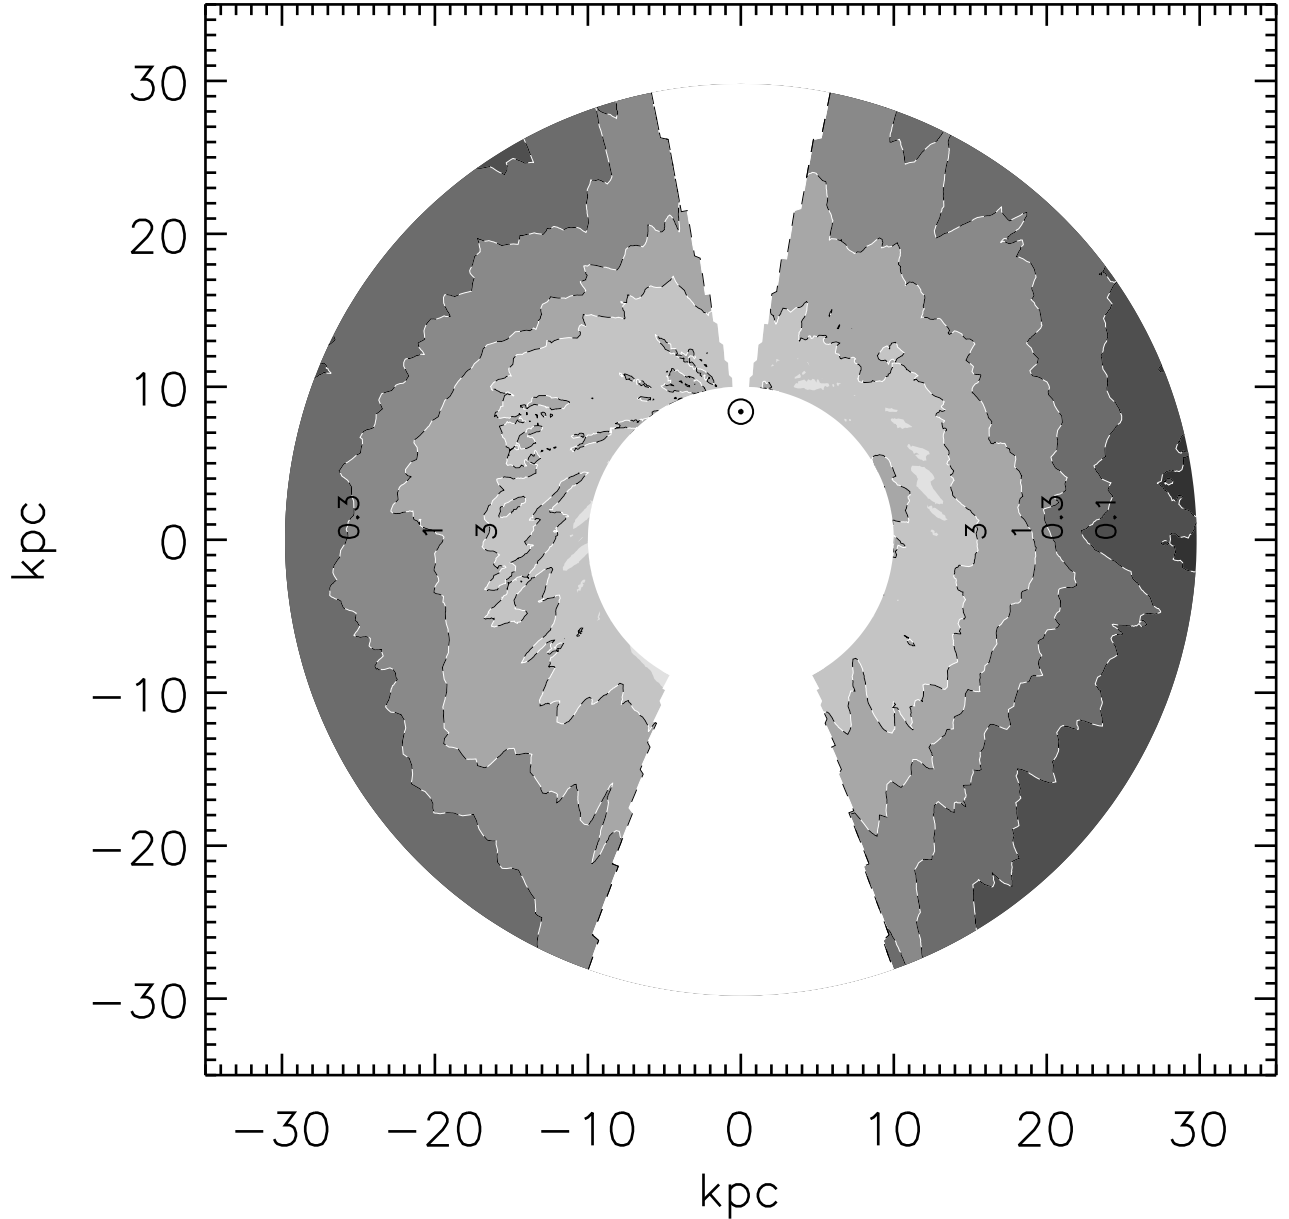
FIG. 1.— A contour map of $\Sigma(R, \phi)$. Contours are evenly spaced in $\log_{10}$ space, at $(3, 1, 0.3, 0.1)$ $M_\odot/pc^2$. The large blank regions directly towards and away from the Galactic center are the excluded regions for which it is hard to establish reliable distances.

plus two Fourier modes with frequency 1 and 2. We fit each ring with the function:

$$W(\phi) = W_0 + W_1 \sin(\phi - \phi_1) + W_2 \sin(2\phi - \phi_2). \quad (11)$$

Each of the three amplitudes W_i and two phases ϕ_i in this fit is a function of radius, because we fit each radius ring independently. We use the gradient-expansion fitting algorithm to perform this fit, weighting each point by the inverse of the squared estimate of the uncertainty in the mean height, $e^2(R, \phi)$. The results of this fit for the rings at $R = 16, 22,$ and 28 kpc are shown in Figures 6,7, and 8. Error bars in these plots represent $e(R, \phi)$.

Each of these plots is a good fit; the offset and the two Fourier components are both necessary and sufficient to describe the large-scale structure of the disk. For illustrative purposes, we will follow the ring at $R = 28$ kpc through each step of the analysis.

Following the evolution of the fit parameters at different places in the disk will tell us how the warp changes with radius. The three amplitude parameters each increase monotonically, with the $m = 0$ mode possibly reaching an asymptotic value near the far end of our radius range (Figure 9). At $R \approx 11$ kpc, the $m = 1$ mode dominates the shape of the warp; the other two
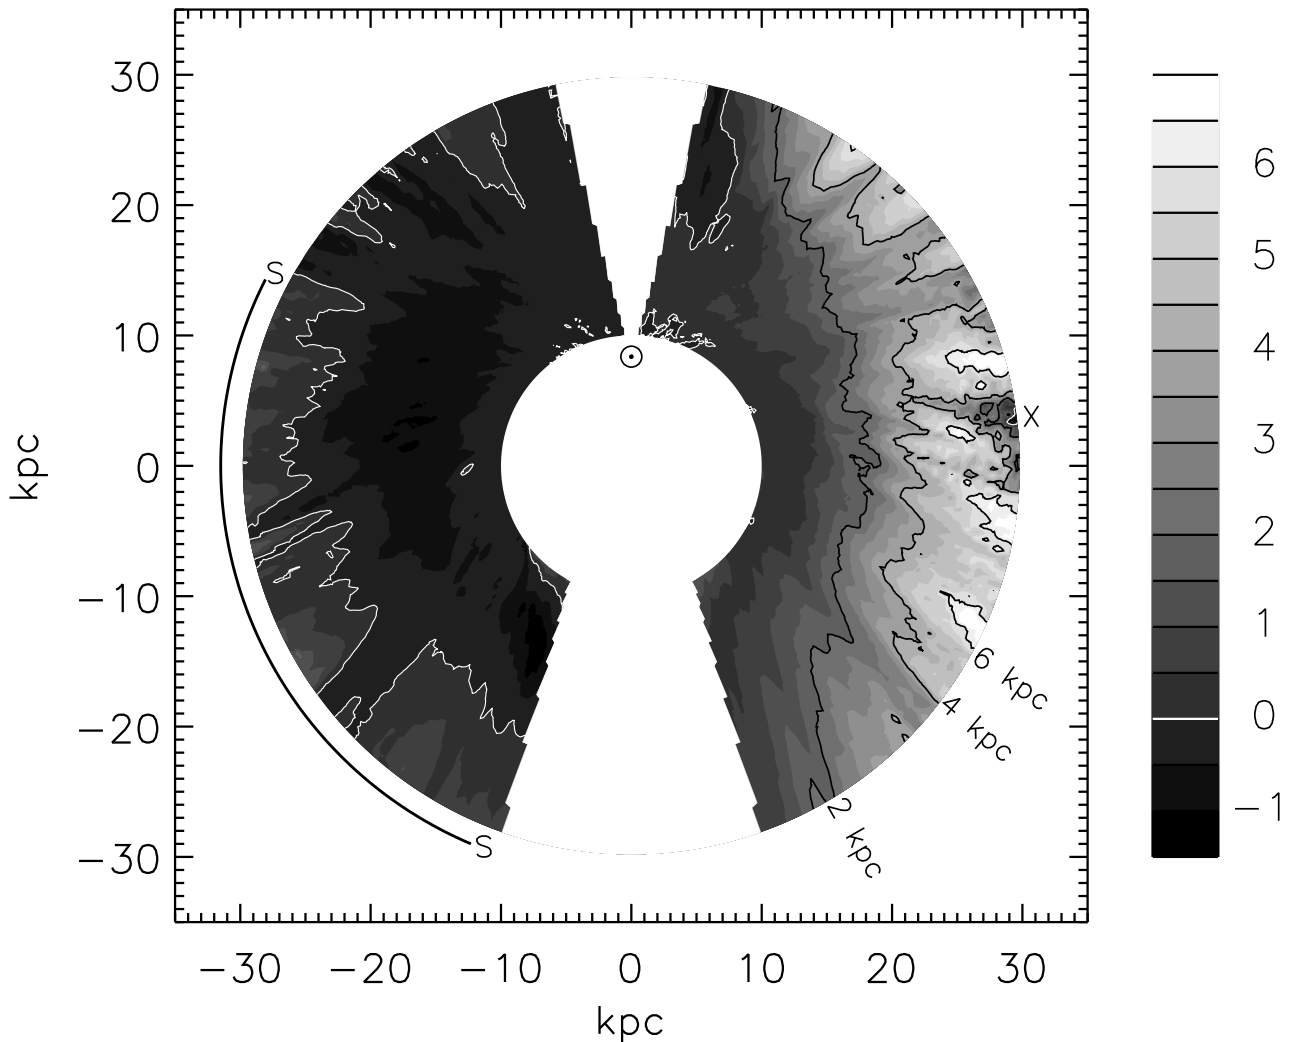
FIG. 2.— A contour map of $h(R, \phi)$. The shading changes every 500 pc in vertical displacement with respect to the $b = 0°$ plane. A white contour is drawn at 0 pc; black contours are drawn at 2000, 4000, and 6000 pc. The colorbar on the right hand side is marked in vertical kpc. The two "S" labels mark the approximate extent of the scalloping as seen by previous authors. The "X" marks the sharp dip discussed in §3.2 and §4.

modes do not become important until $R \approx 15$ kpc. This plot implies that the $m = 1$ mode has power even at the edge of our grid, thus we cannot establish the onset of the warp. We do a linear least-squares fit on the growth of each warp parameter using the function

$$W_n = k_0 + k_1 (R - R_k) + k_2 (R - R_k)^2 \qquad (12)$$

where only points at R_k and beyond are weighted in the fit. R_k is arbitrarily chosen to be near where each of the three modes starts growing. The value of k_0 for each fit is strongly correlated to the choice of R_k.

Binney & Merrifield (1998) (hereafter BM) discuss an approximation to the warp that is of similar form to our fit. They also fit the first three modes, but they fix the line of zeros to lie along the Sun-Galactic center line. In our fit, this would be equivalent to setting ϕ_1 and ϕ_2 to zero. Also, BM fix W_0 and W_2 to be the same. Note that our warp data are adjusted for elliptical gas orbits, and

TABLE 1
PARAMETERS RESULTING FROM A LINEAR LEAST-SQUARES FIT
TO THE WARP

m	R_k (kpc)	k_0 (pc)	k_1 (pc kpc^{-1})	k_2 (pc kpc^{-2})
0	15	-66	150	-0.47
1	10	9	197	-3.1
2	15	-70	171	-5.3

are filtered according to the procedure described in §2, whereas BM's data are not. Fig. 10 compares the mode amplitudes calculated from the data, the fit from BM, and our fit. BM overestimate the strength of the warp starting at around $R \approx 22$ kpc for the $m = 2$ mode, 24 kpc for the $m = 0$, and 27 kpc for the $m = 1$. However, the BM fit matches our data fairly well for the radii where
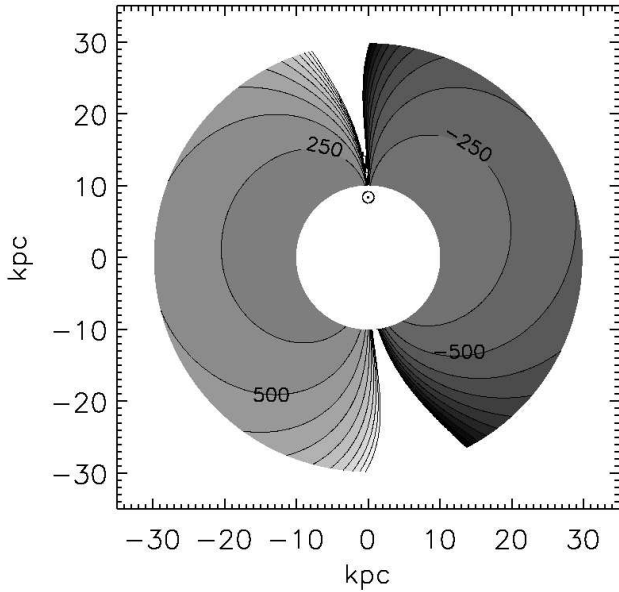
FIG. 3.— A contour map of dr/dv_r, including the velocity contribution from the elliptical gas orbit correction. The axes are marked in kpc; the contours delineate regions of 250 pc/(km s^{-1}). The contours rise close to the Galactic center and anticenter due to the small radial velocities of gas in those regions, demonstrating why the data in those directions are excluded from the rest of the study.

the warp is growing most rapidly.

The line of maxima for the $m = 1$ mode and one of the two lines of maxima for the $m = 2$ mode are marked on the contour map of the warp fit (Figure 11). Since ϕ_1 and ϕ_2 are actually the line of zeros, the lines of maxima are shifted 90° and 45° from these values in our fit, respectively. The parameter ϕ_2 is not well determined at small radii in our fit because the amplitude of the $m = 2$ Fourier component is very small in that region. For this reason, we do not plot ϕ_2 in the region where W_2 is less than 150 pc. There is little evidence for precession in the lines of maxima for the two modes, and the line of maxima for the $m = 1$ mode is roughly aligned with one of the lines of maxima of the $m = 2$ mode; for each radius their difference in ϕ is less than 12°. Burton (1988) has examined how the two lines of zeros for the mean height change in ϕ as a function of radius; they appear to stay roughly aligned with the Sun-Galactic center line as R changes from R_0 to 26 kpc. Fig. 11 shows that the line of zeros for the sum of our three warp modes falls within the excluded region, consistent with the earlier work.

The three component fit shown in Figure 11 does a good job of reproducing the large-scale features in the mean height map. This is not surprising, given that the $m = 0, 1$, and 2 modes are the strongest in the Lomb periodogram; differences between the two plots are due to power in high frequency modes. We will now examine the differences between the mean height map and the warp fit.

3.1.2. *Scalloping*

From the fit to the warp, we can determine a function for the scalloping: $s(R, \phi) = h(R, \phi) - W(R, \phi)$. We

continue to assume no information regarding the shape of the disk in the excluded regions. We can now look for higher m modes of oscillation that may be present in $s(R, \phi)$. Such a signal would be present if the gas were coherently moving up and down as a function of ϕ. The Lomb periodogram is a useful numerical method for detecting periodic signals in unevenly spaced data (Press et al. 1992). It provides the same amplitude we would determine by using a linear least squares fit, even if the signal is not sinusoidal, while allowing for a straightforward error analysis. The data are weighted equally per point, which is necessary to deal with unevenly sampled data. To ease comparison with the warp component amplitudes in Fig. 9, we use an unnormalized Lomb periodogram and take the square root of the power to get the amplitude.

We use a Monte Carlo algorithm to determine the noise level for the Lomb periodogram. For each ring, the null hypothesis is Gaussian white noise with the same dispersion as the data in $s(R)$. We construct 10^3 sets of noise for each ring, run the Lomb periodogram on each set, and record the highest peak in the Lomb amplitude. We then determine the distribution of peaks in these amplitudes, and define the 95% confidence interval as the amplitude just larger than 95% of the noise amplitude peaks. We use the same technique to determine the 99% confidence level. To conclude that a signal is real and not caused by noise, the signal strength must cross these confidence levels. We found nearly identical noise levels by scrambling the order of the data instead of using Gaussian white noise, indicating that these calculations are robust.

The Lomb periodogram of $s(R, \phi)$ for the $R = 28$ kpc ring is shown in Figure 12. 95% and 99% confidence levels are marked as thresholds in amplitude. At this radius, there is significant strength in modes 4-6, 10, and 15. These modes are markedly weaker than the warp components at the same radii, which have amplitudes of 1-2 kpc.

In the case of unevenly spaced or missing data points, simple sine waves are not eigenmodes of the system. One way to see this is to take a Lomb periodogram of a pure sine wave on our ϕ grid with missing data; a small amount of the power will leak into other frequencies. Thus subtracting out the warp from $h(R, \phi)$ influences the amplitudes of the higher order modes from the Lomb periodogram because individual m modes are not independent. We argue that the subtraction is nonetheless acceptable because the warp and the scalloping appear to be due to physically distinct phenomena; studying them is much easier once they have been separated. Removing the strongest modes will also result in a more accurate power spectrum of the weaker frequencies, since we eliminate the power leakage from the stronger to the weaker modes.

Figure 13 shows how the amplitudes in several different m modes evolve with R. Modes 4-6, 10, and 15 appear to increase in strength at outer radii, while mode 3 is strongest in the intermediate radii in our map. Other modes have no detectable strength because they do not cross the 95% confidence level at any radii. Also notice how the confidence thresholds rise with radius because of the increase in the dispersion of $s(R, \phi)$.

Another consequence of using the Lomb periodogram

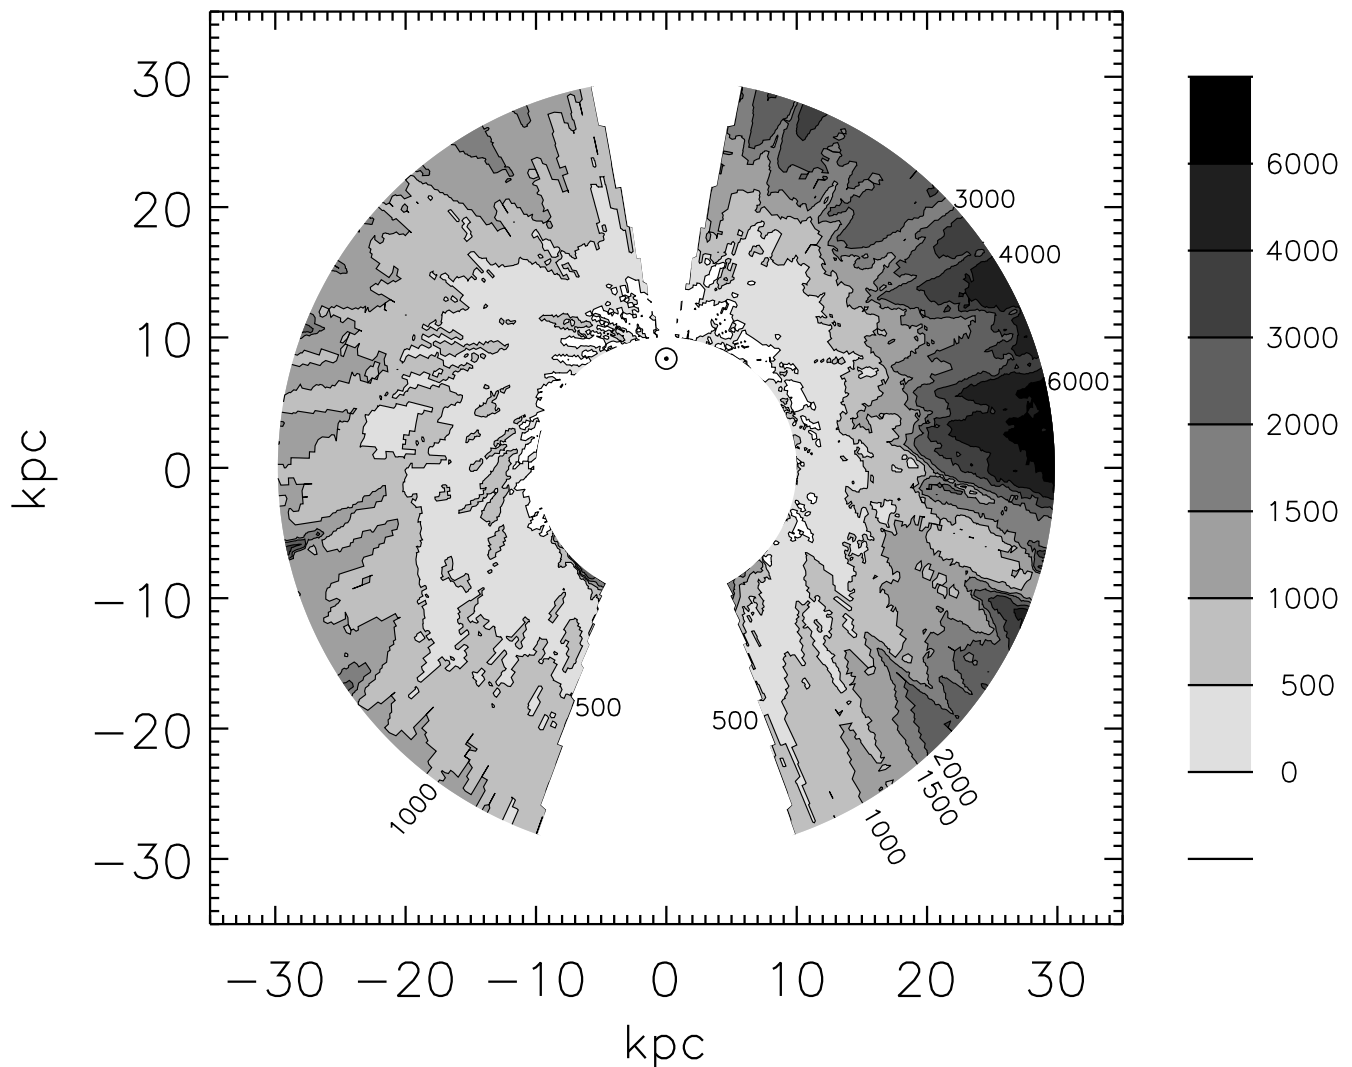FIG. 4.— A contour map of $T_h(R, \phi)$; the thickness of the disk is approximately twice this quantity. Contours are drawn at 0, 500, 1000, 1500, 2000, 3000, 4000, and 6000 pc. The colorbar and the figure labels are in pc.

analysis on unevenly spaced data is that it is possible to be fooled by a false signal due to ringing from interference between different m modes. It is difficult to protect against this, but if two strong modes were interfering with each other to produce a third signal, we would expect at least two of the signals to increase in strength at the same radius. The strongest modes in Figure 13 become significant at different radii and have visually different evolution with radius; we conclude that these modes are real.

3.2. *Local Structure*

The Lomb periodogram cannot be used to study the local structure of the disk since it cannot determine where in azimuth each mode is strong. Imagine a situation where, like a falling stone creating ripples on the surface of a pond, something passes through the HI disk and excites a local series of vertical oscillations. The oscillation will add power to some frequencies in the Lomb

periodogram, but this effect may be dwarfed by oscillations elsewhere in the disk. We wish to detect these sorts of perturbations and learn where in the disk they are prominent.

One way to draw out this type of structure is to use a wavelet analysis. Wavelets are ideal for our purposes because they are localized in both position and frequency space; in other words, they will show where in the disk a given mode of oscillation is dominant. In wavelet theory, it is beneficial to choose a mother wavelet that is similar in shape to the signals you are searching for. Since we are looking for sinusoidal perturbations, we use the normalized Morlet wavelet with $\omega_0 = 6$; this is simply a sine wave times a Gaussian envelope. In the same way that a Fourier transform breaks down the frequency structure of a signal using sines and cosines, a wavelet transform breaks down a signal in terms of a set of these Morlet wavelets centered at different spatial positions and with a range of frequencies. However, in a continuous wavelet
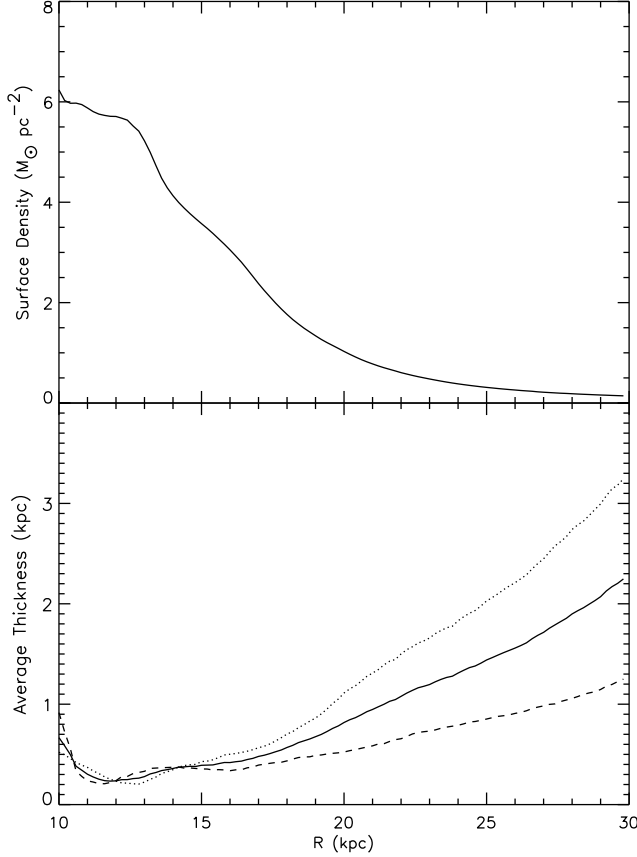
FIG. 5.— The upper panel is the disk surface density, and the lower is the half thickness of the gas as a function of radius. In the lower panel, the solid line is the average thickness over the whole disk, the dotted line is the average from $0° < \phi < 180°$, and the dashed line is from $180° < \phi < 360°$.
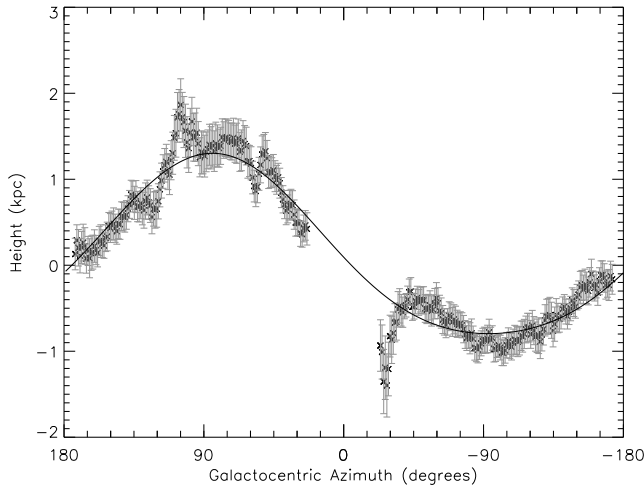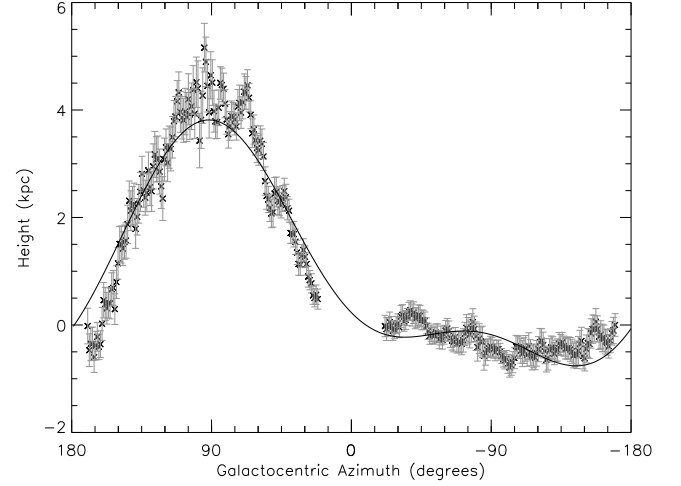

FIG. 7.— The solid line is the warp fit $W(\phi)$ to h at $R = 22$ kpc. The approximate error in the mean height of the disk at each point is represented by the error bars.
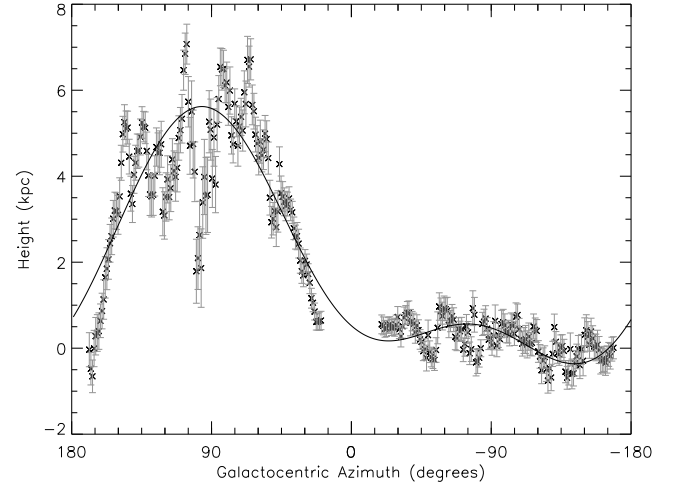

FIG. 8.— The solid line is the warp fit $W(\phi)$ to h at $R = 28$ kpc. The approximate error in the mean height of the disk at each point is represented by the error bars.

FIG. 6.— The solid line is the warp fit $W(\phi)$ to h at $R = 16$ kpc. The approximate error in the mean height of the disk at each point is represented by the error bars.

transform like the one we perform, the different wavelet functions are not orthogonal.

To avoid having to interpolate or zero out the excluded regions, we will examine the northern and southern halves of the Galaxy separately. As in §3.1.2, we will work with the mean height function once the warp has been subtracted, $s(R, \phi)$. Points in frequency and position space that would involve the excluded regions are defined as being inside the "cone of influence"; points in this region are subject to edge effects, and are therefore discarded. The points affected by this are not just those in the excluded regions but also those adjacent to these regions, within some range set by the wavelength (and thus the frequency) of the particular mode (Torrence & Compo 1998). Thus, lower frequency modes will have a larger portion of the disk fall inside the cone of influence, and have to discard a larger range of points.

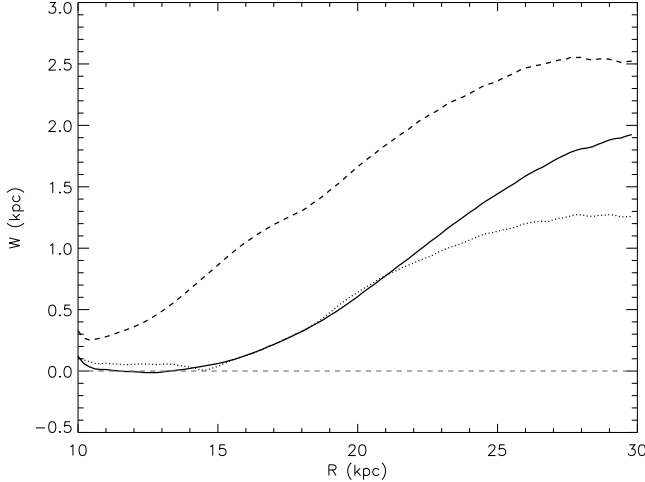It is also important to have some analytic measure

FIG. 9.— The evolution of the warp amplitude parameters as a function of radius. W_0 is the solid line, W_1 is the dashed, and W_2 is the dotted. The dashed grey line is zero amplitude.
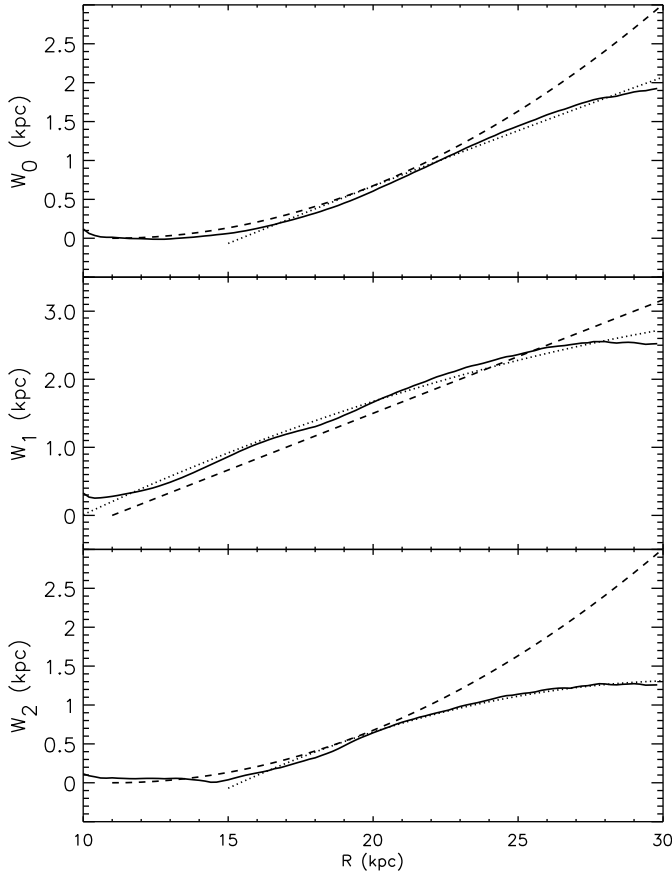

FIG. 10.— In each panel, the solid line is the calculated amplitude of the the mode from (11), the dashed line is the fit from BM, and the dotted line is the fit described in (12). Top panel: $m = 0$. Middle panel: $m = 1$. Bottom panel: $m = 2$.

of which peaks in the filtered power spectrum are significant. Significance levels are discussed in detail in Torrence & Compo (1998). We construct a combination of parameters that have a χ^2 distribution, and count

as significant those that cross the 95% confidence level threshold. Again, we use Gaussian white noise to model the randomness in the height function.

The wavelet power spectrum, $\overline{W_n}^2$, for the northern half of the $R = 28$ kpc ring is plotted in Figure 14; the southern half is shown in Figure 15. The mathematical details of the wavelet transform are summarized in the Appendix. This ring has several regions with significant power. These regions have a large width in both dimensions because wavelets do not have precise resolution in either position or frequency space. Much like the uncertainty principle, the cost of using a technique that gives both position and frequency information is mediocre resolution in both. The power around $\phi \approx 90°$ comes from the sharp dip at that azimuth; as the figure shows, this causes ringing for a large range of frequencies. The power at $-45° \gtrsim \phi \gtrsim -90°$ is due to the "scalloping" in that region.

The filtered wavelet power spectra for several different bands are plotted in Figure 16. The band that is labeled $m = 7$ is actually the filtered sum of the power spectra that satisfy $6.5 \leq m \leq 7.5$, and so on. Note that since we calculate the power spectra for the northern and southern halves independently, the significance contours are different for the two hemispheres, in addition to being a function of R. In practice, this occurs because the variance of $s(\phi)$ for the northern half of the Galaxy is larger than that for the southern, and the strength required of the power spectrum to cross the 95% confidence level is directly proportional to the variance.

4. DISCUSSION

We grouped the $m = 0, 1$ and 2 modes together in the warp because of their similar magnitudes in a Lomb periodogram. A close look at their dependence with R (Fig. 9) demonstrates that there are additional similarities. Both the 0 and the 2 mode are near zero until about 15 kpc from the Galactic Center. The 0 mode grows linearly from this point outwards while the 2 mode declines a small amount, and then grows approximately linearly by about 1.2 kpc. In comparison, the $m = 1$ mode starts out fairly large (≈ 300 pc), declines slightly and then grows by about an order of magnitude. Though the $m = 2$ mode does break off from this pattern at larger radii, all three components grow essentially monotonically, approximately linearly, and with similar slopes over the range $15 \lesssim R \lesssim 22$ kpc. This may be a clue that their origins involve the same physics, and helps to justify our classification of higher frequency modes as scalloping. Studying the radial dependence of the warp inside the solar circle seems a fruitful way to learn more about these three components.

The filtered wavelet power spectrum maps (Fig. 16) are a representation of the scalloping in the outer Galaxy. Due to frequency-position uncertainty relations, it is impossible to establish precisely the oscillation frequency of any local disturbance. This uncertainty manifests itself in the wavelet transform by perturbations that are only somewhat localized in frequency and position space, and therefore have some width in both.

These maps demonstrate that the $m \approx 10$ scalloping found near $\ell \approx 310°$ and $R \approx 25$ kpc mentioned in previous work (Henderson et al. 1982; Kulkarni et al. 1982) is real. The strong power around $m = 10$ in the wavelet
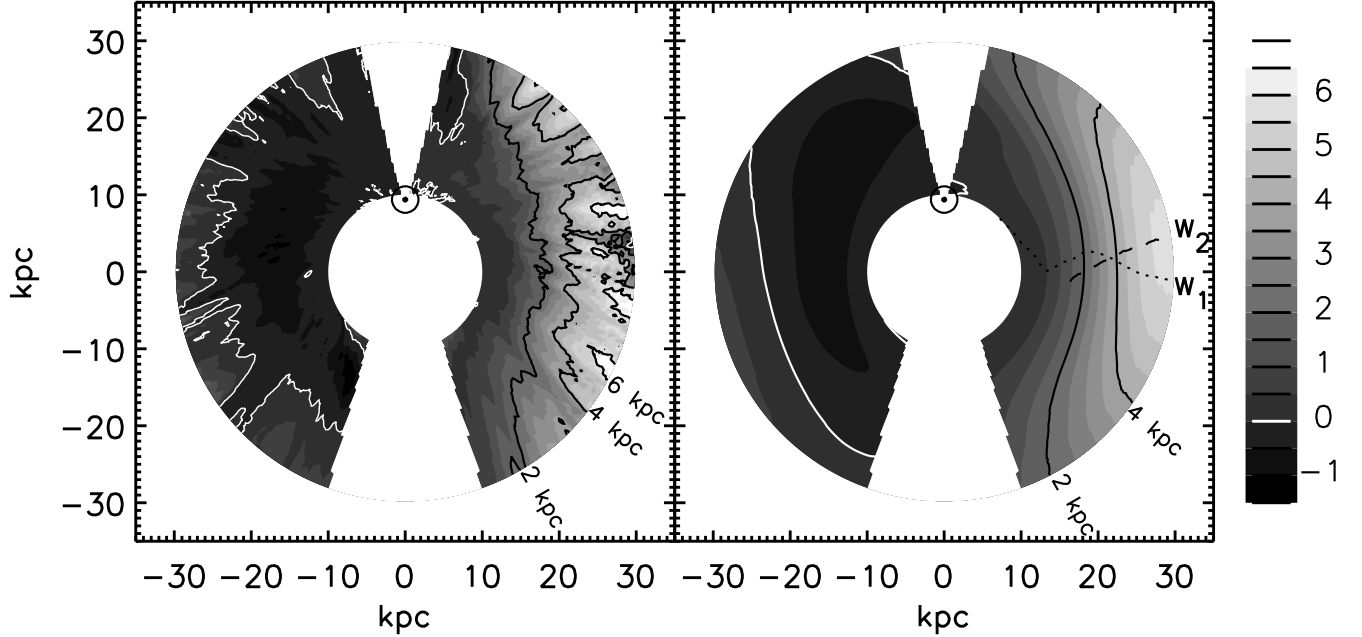
FIG. 11.— Left panel: $h(R, \phi)$. Right panel: The fit to the warp is plotted, along with the lines of maximum amplitude for the $m = 1$ (dotted) and 2 (dashed) modes. These lines are marked W_1 and W_2 respectively. The white contour line denotes a height of 0 pc; black lines mark the 2000, 4000, and 6000 pc elevations.
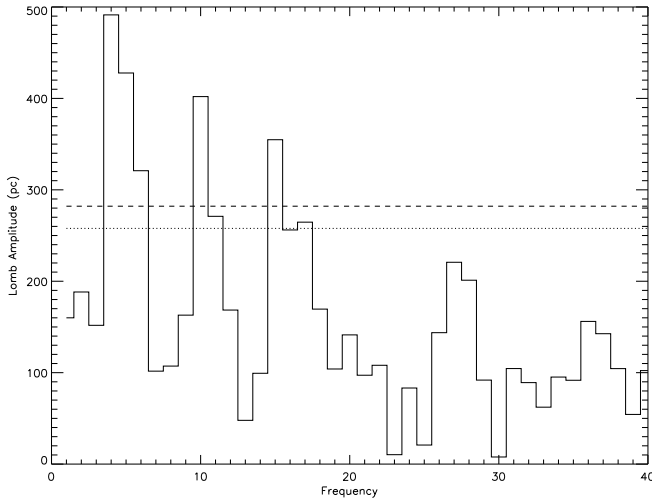

FIG. 12.— The Lomb Periodogram for the ring at 28 kpc, once the warp has been removed. The 95% and 99% significance levels are drawn as dotted and dashed lines, respectively.

transform is accompanied by a large amplitude in the Lomb periodogram for $m = 10$ at the corresponding radii.

Wavelet transforms are also subject to the same problems that bedevil traditional Fourier transform approaches. For example, there is a large amount of power around $R \approx 30$ kpc and $\phi \approx 90°$ in $8 \leq m \leq 15$. This power is most likely not a result of scalloping in all of these different frequencies over some range of ϕ; instead, it is probably due to ringing. As in a Fourier decompo-

sition, sharp changes in height will cause ringing in all frequencies of a wavelet transform. Indeed, the height map (Fig. 2) does have an abrupt drop near $\phi \approx 90°$ which could cause this ringing; we mark this region with an 'X'. The same feature can be seen after the warp has been subtracted in the ring at 28 kpc plotted in Figure 14.

We use the local and global analyses in conjunction to determine whether there is any scalloping mode that is present over a full 2π ring in the outer Galaxy. This is important because it will determine whether the mechanism that causes scalloping operates on a global or local scale. For this task, both techniques are necessary because even a pure $m = 10$ oscillation will have some width in frequency space when put through a wavelet transform. However, this same oscillation will have a sharp peak in the Lomb periodogram. Therefore, in order to state that some mode exists over a full 2π, we require both significant strength in the Lomb periodogram and significant power over a large range of ϕ in the wavelet transform. This combination will conclusively determine whether the scalloping is a global or local phenomenon.

For modes with $m \leq 6$ this technique is not useful. These modes are clearly important because their Lomb periodograms show significant power for $3 \leq m \leq 6$. Unfortunately, these modes have large enough wavelengths that the excluded regions always interfere with the wavelet transform regardless of where in the disk we look. A method of reducing the size of the excluded regions would alleviate this problem. However, this is likely to be difficult because eliminating the excluded regions would require a detailed knowledge of both v_R and the distribution of gas velocities due to turbulence in the
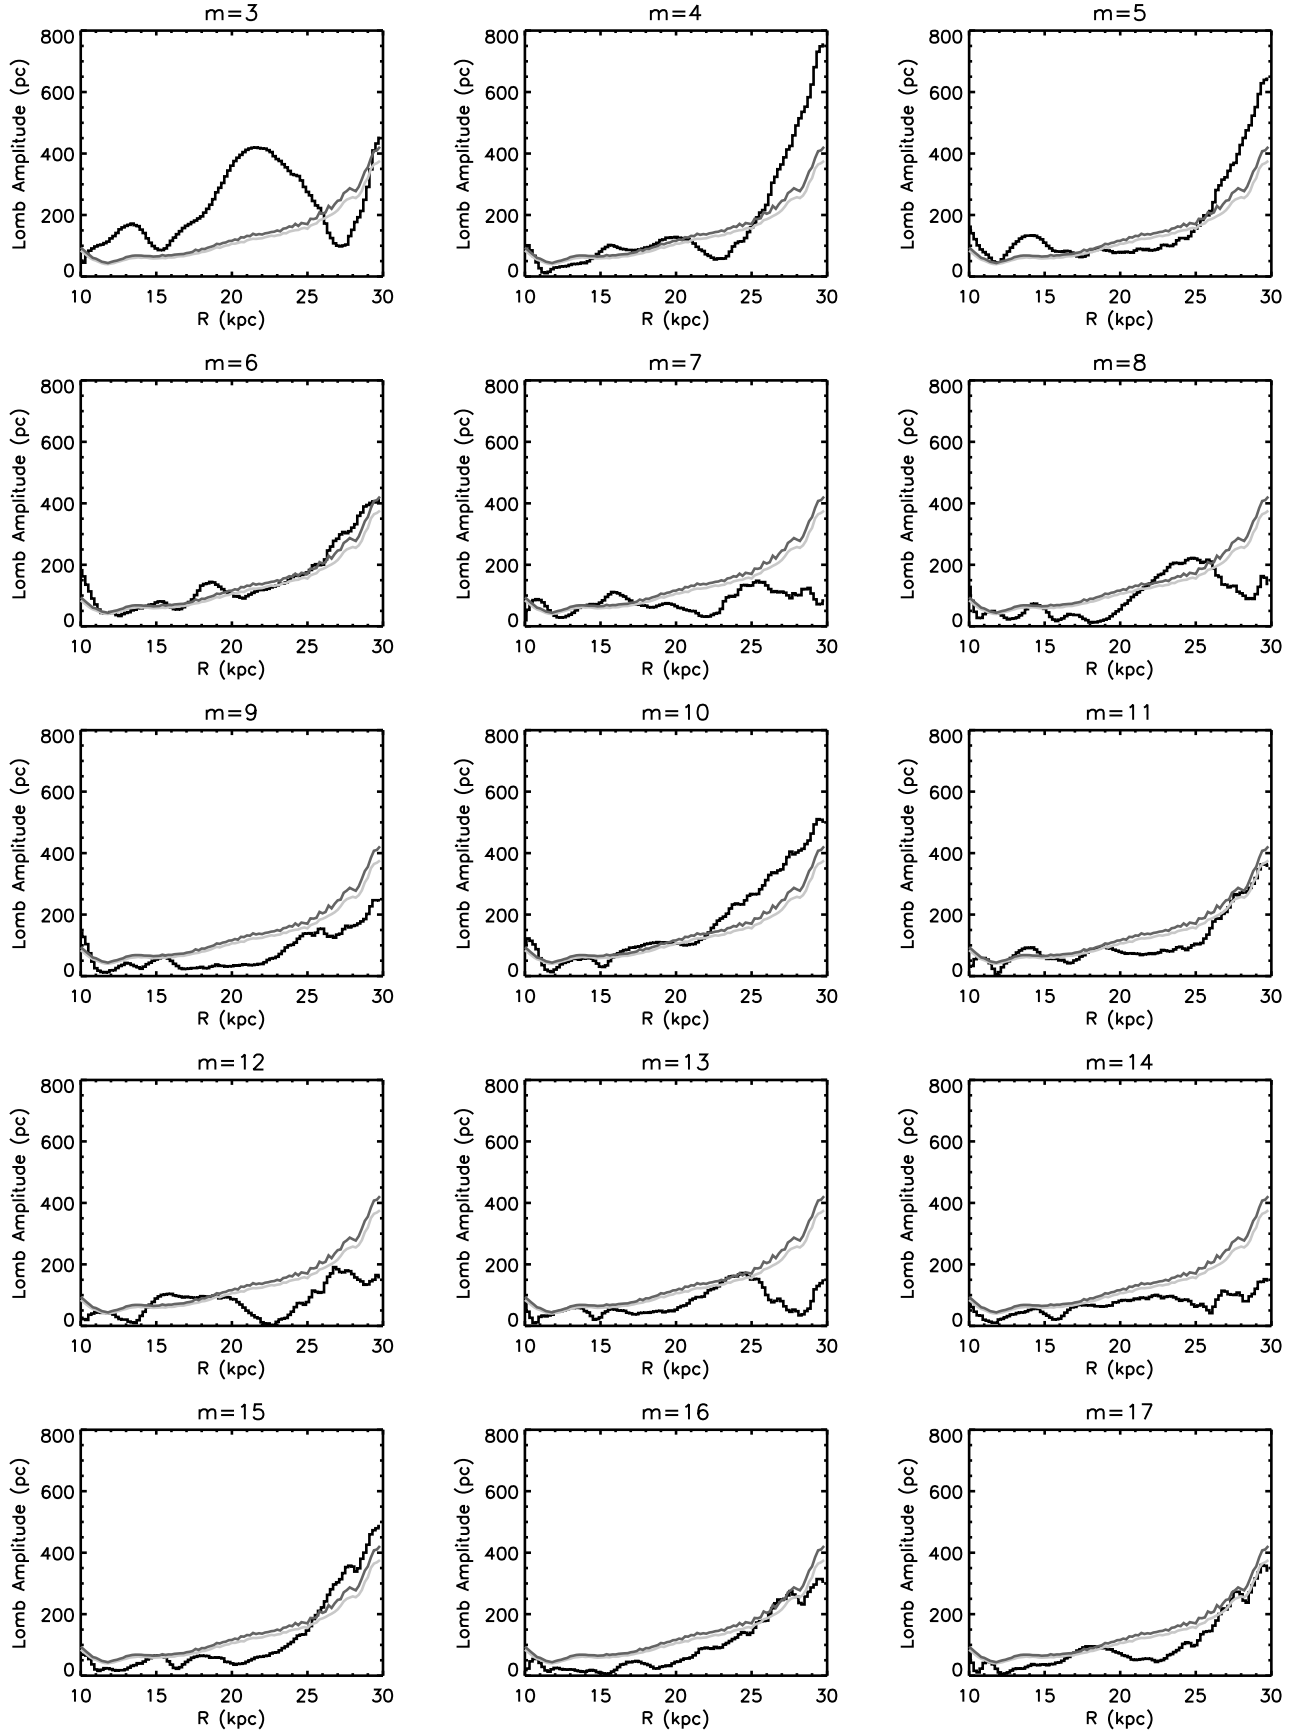
FIG. 13.— The Lomb periodogram amplitudes as a function of radius, once the warp has been removed. The 95% and 99% significance levels are drawn as light grey and dark grey lines, respectively.
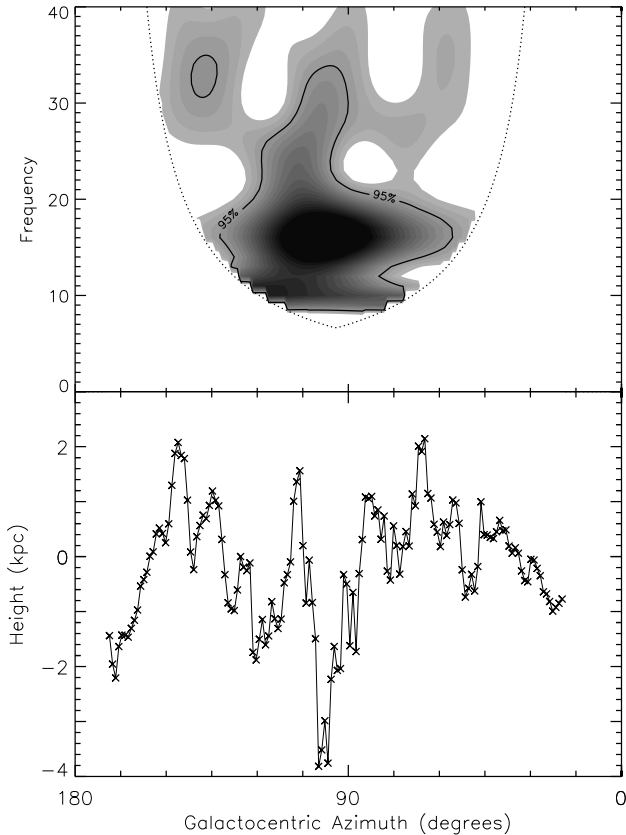
FIG. 14.— The filtered wavelet power spectrum for the northern half of the $R = 28$ kpc ring once the warp has been subtracted. The 95% significance level is marked with a solid line, and the cone of influence is outlined with a dotted line. The function $s(28 \text{ kpc}, \phi)$ is plotted in the lower panel.
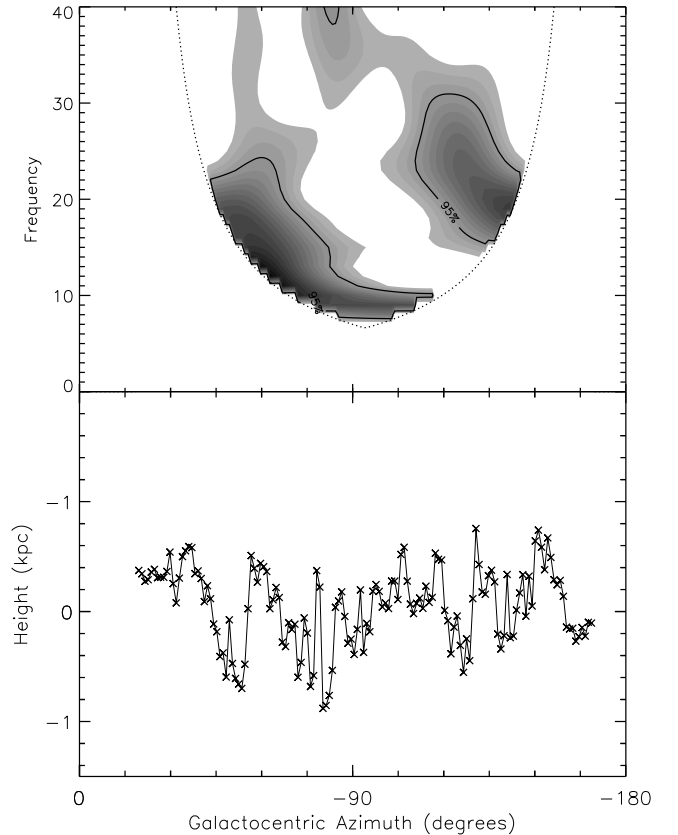
FIG. 15.— The filtered wavelet power spectrum for the southern half of the $R = 28$ kpc ring once the warp has been subtracted. The 95% significance level is marked with a solid line, and the cone of influence is outlined with a dotted line. The function $s(28 \text{ kpc}, \phi)$ is plotted in the lower panel.

disk. Portions of these region are optically thick, which will make density determination impossible.

With these constraints, the Lomb periodogram leads us to conclude that the modes $m = 10$ and 15 are the most fruitful places to look for scalloping that travels a full circle around the disk. Other modes do have interesting features, but their smaller Lomb amplitudes imply that we could be fooled by the imperfect frequency resolution of the wavelet power spectrum. For example, the modes $7 \leq m \leq 9$ have significant wavelet power for large portions of the northern half of the galaxy, but none of these modes has significant signal for the corresponding radius range in the Lomb periodogram.

For the $m = 10$ mode the only large region with significant power in the south is the one around $\ell \approx 310°$ that we noted previously. The power in this perturbation has fallen below the 95% confidence level by $\ell \approx 270°$, implying that this scalloping is a local effect.

The same appears to be true for the $m = 15$ oscillations, although this mode does have a region of high significance near $R \approx 30$ kpc over a large range of ϕ. However, the northern part of this signal is the region where we believe ringing to play an important role. It is therefore possible that the $m = 15$ mode does have significant wavelet power over the full 2π near $R \approx 30$ kpc, but the evidence is not conclusive because the northern part of the wavelet power is most likely not due to scalloping.

With the exception of the $m = 15$ mode, no frequency we examine with the wavelet transform carries significant power around an entire ring along with a correspondingly significant Lomb periodogram strength. For this reason, we conclude that scalloping generally appears to be a local phenomenon. The modes $3 \leq m \leq 6$ remain a possible exception, since we were not able to study them with a wavelet transform.

It remains unclear what mechanism acts as the energy source for the scalloping. One possible cause is a massive object passing through the disk that excites local vertical oscillations in the HI gas. Azimuthally traveling wavefronts can be created by the magnetic field threaded through the disk; if the field is primarily azimuthal in nature, vertical oscillations will have larger phase and group velocities in the azimuthal direction than in the radial direction. This could lead to scalloping such as that seen towards $\ell \approx 310°$.

5. CONCLUSIONS

We fit the global shape of the warp on our grid of concentric rings using a vertical offset and two sinusoidal modes. Outside of $R \approx 20$ kpc, each of these three modes has more power than any of the higher frequencies we look at. The amplitude increases with radius over our entire radius range for the 1 mode, and starting from around 15 kpc for the 0 and 2 modes. The growth of the 0 and 2 modes results in asymmetry in the warp; this
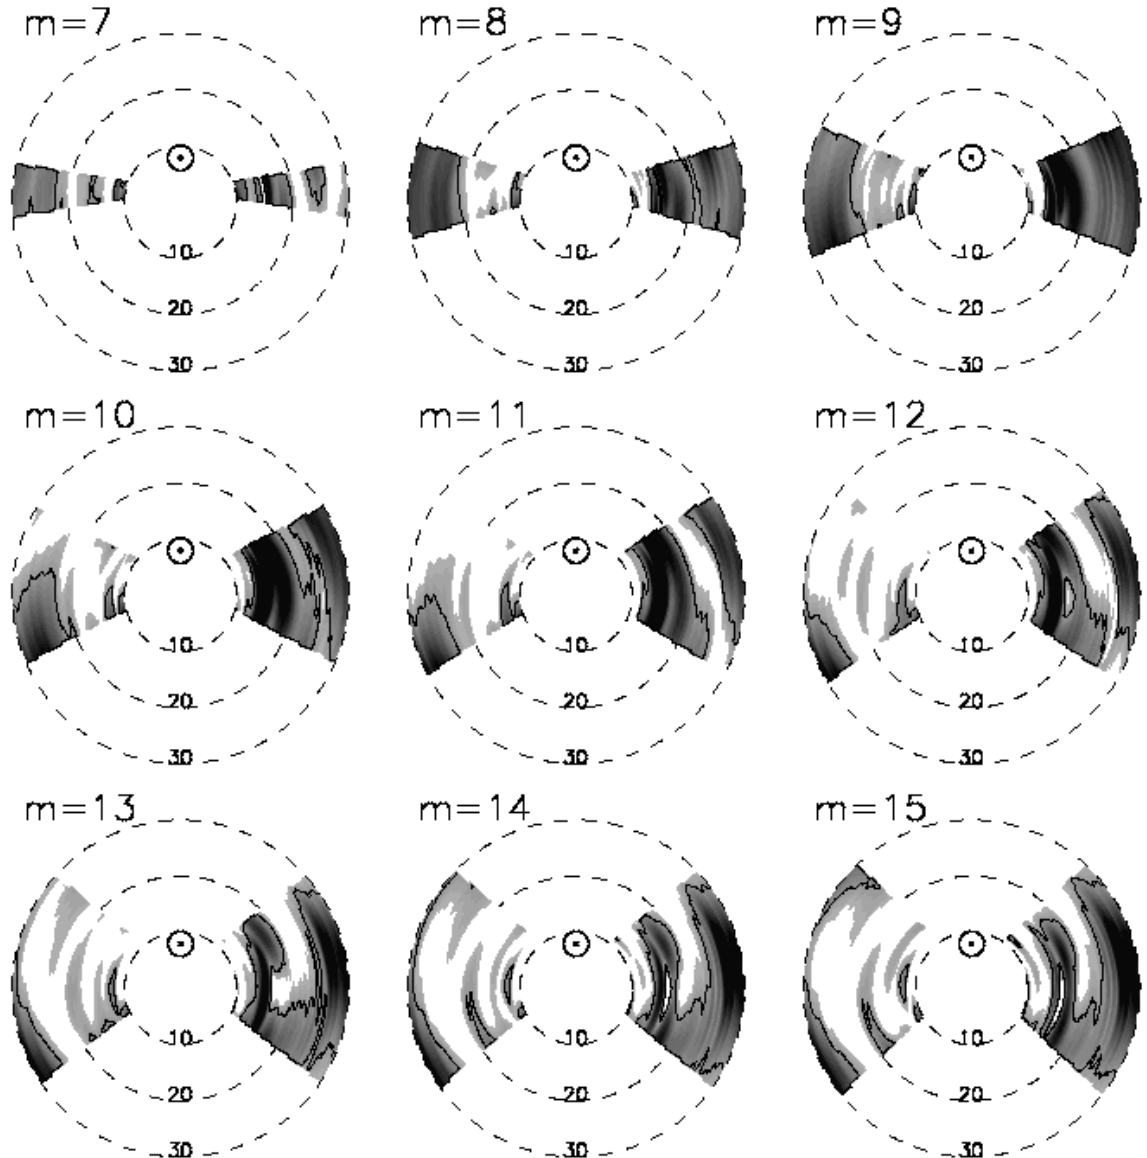
FIG. 16.— The filtered wavelet power spectrum for the outer Galaxy. Darker contours represent higher power, and the 95% significance level is marked with a solid line. The labels on the rings give the distance from the Galactic center in kpc.

growth begins near where the stellar disk ends. The line of maxima of the $m = 1$ mode is essentially coincident with one of the lines of maxima of the $m = 2$ mode. There is little evidence for precession or winding of these two modes.

A global analysis with the Lomb periodogram shows that each m mode evolves differently with radius. The most interesting include $m = 3 - 6, 10$, and 15, each of which start at small radii below the 95% significance level and then cross it further out in the disk. An analysis combining the global Lomb periodogram and the local wavelet transform shows that none of the modes with $7 \leq m \leq 15$ have strength over a full ring of the disk.

Using a wavelet transform, we show that the scalloping observed by previous authors near $\ell \sim 310°$ is real. We therefore conclude that the scalloping is a local effect. Lower frequency modes proved impossible to study with wavelets due to the presence of the excluded regions.

APPENDIX

MEASURING RADIAL MOTION OF THE GAS

This appendix describes our method for finding a functional form for $v_\Pi(R)$, the magnitude of the elliptical corrections to circular rotation. A naive surface density plot using only circular rotation has a large degree of asymmetry across the lines $\ell = 0°$ and $180°$ (Kulkarni et al. 1982). Figure A17 shows the calculated surface density without any correction for elliptical gas orbits, but with the filters we describe in §2. Observational experience tells us that the surface density should not have discontinuities in cardinal directions; the elliptical orbits we describe in §2 provide the strongest corrections in the directions where the $\cos\phi$ term in (1) is large, and smaller corrections elsewhere. The algorithm is based on matching the surface density of the HI disk on either side of the excluded region centered at $\ell = 180°$. We draw confidence from the fact that this fit also does a good job matching the contours around $l = 0°$ even though they are not included in the fit.

In order to calculate a surface density, we must first know $v_\Pi(R)$ because it enters into the derivative $|dv_r/dr|$. Thus, determining $v_\Pi(R)$ from anything connected to ρ is a circular problem. We must assume some form of $v_\Pi(R)$ and check to see if it results in surface densities that are well matched. We choose the functional form:

$$v_\Pi(R) = \alpha\frac{(R - R_0)}{R_0} + \beta\frac{(R - R_0)^2}{R_0^2}, \tag{A1}$$

and do not include a zeroth order term to ensure that $v_\Pi(R)$ passes through zero for the solar circle. Radhakrishnan & Sarma (1980) showed that the LSR does not have a radial velocity with respect to the Galactic center.

Our algorithm follows these steps:

1. Construct $v_\Pi(R)$ for some combination of α and β

2. Interpolate from the LAB survey to find $T_b(R, \ell, z)$ for a grid in R_i and z_j over the survey range $|b| \leq 30°$ and $155° \leq \ell \leq 165°$ or $195° \leq \ell \leq 205°$ using $v_\Pi(R)$

3. Sum over z to find the surface density $\Sigma(R, \ell)$

4. Average over ℓ for the two subsets of ℓ to find $\Sigma_{165}(R)$ and $\Sigma_{195}(R)$

5. Calculate a modified χ^2 statistic to determine how well the surface densities are matched

The modified χ^2 is defined as

$$\chi^2 = \sum_{R_i}\left[\frac{\Sigma_{165}(R_i) - \Sigma_{195}(R_i)}{\Sigma_{165}(R_i) < \Sigma_{195}(R_i)}\right]^2 \tag{A2}$$

where the $<$ operator returns the smaller of its two operands.

Using this algorithm we search for the values of α and β that minimize the modified χ^2. We find good matches for $\alpha = 8.67$ and $\beta = -1.08$.

We also tried a weighting that only fit points beyond the perturbations of the spiral arms, i.e. $R > 2R_0$. Although this did change $v_\Pi(R)$ by 25% or so, it had no qualitative effect on our other results.

WAVELETS

The continuous wavelet transform for a discrete series of points z_n is given by

$$W_n(s) = \sum_{n'=0}^{N-1} z_{n'}\psi^*\left[\frac{(n' - n)\Delta\phi}{s}\right] \tag{B1}$$

where N is the number of points in the series, n is a position index, $\Delta\phi$ is the spacing of the points in ϕ space, and ψ^* is the complex conjugate of the normalized Morlet wavelet (Torrence & Compo 1998):

$$\psi(\eta) = \left(\frac{\Delta\phi}{s}\right)^{1/2}\pi^{-1/4}e^{i\omega_0\eta}e^{-\eta^2/2}. \tag{B2}$$

The scale of the transform is s; for the Morlet wavelet this is simply related to the wavelength by $\lambda = 1.03s$. Evenly spaced points are necessary in order to use this transform, but since we are treating the two halves of the Galaxy separately, $s(R, \phi)$ is split into two halves, each with equally spaced points. We perform this transform for the dense set of frequencies given by

$$s_j = s_0 2^{j\Delta j}, \quad j = 0, 1, \ldots, J. \tag{B3}$$

Here, s_0 is the smallest scale that can be sampled, $2\Delta\phi$. Δj is a measure of how densely we sample in scale space; because computation time is not large, we choose relatively dense sampling throughout: $\Delta j = 0.0125$.
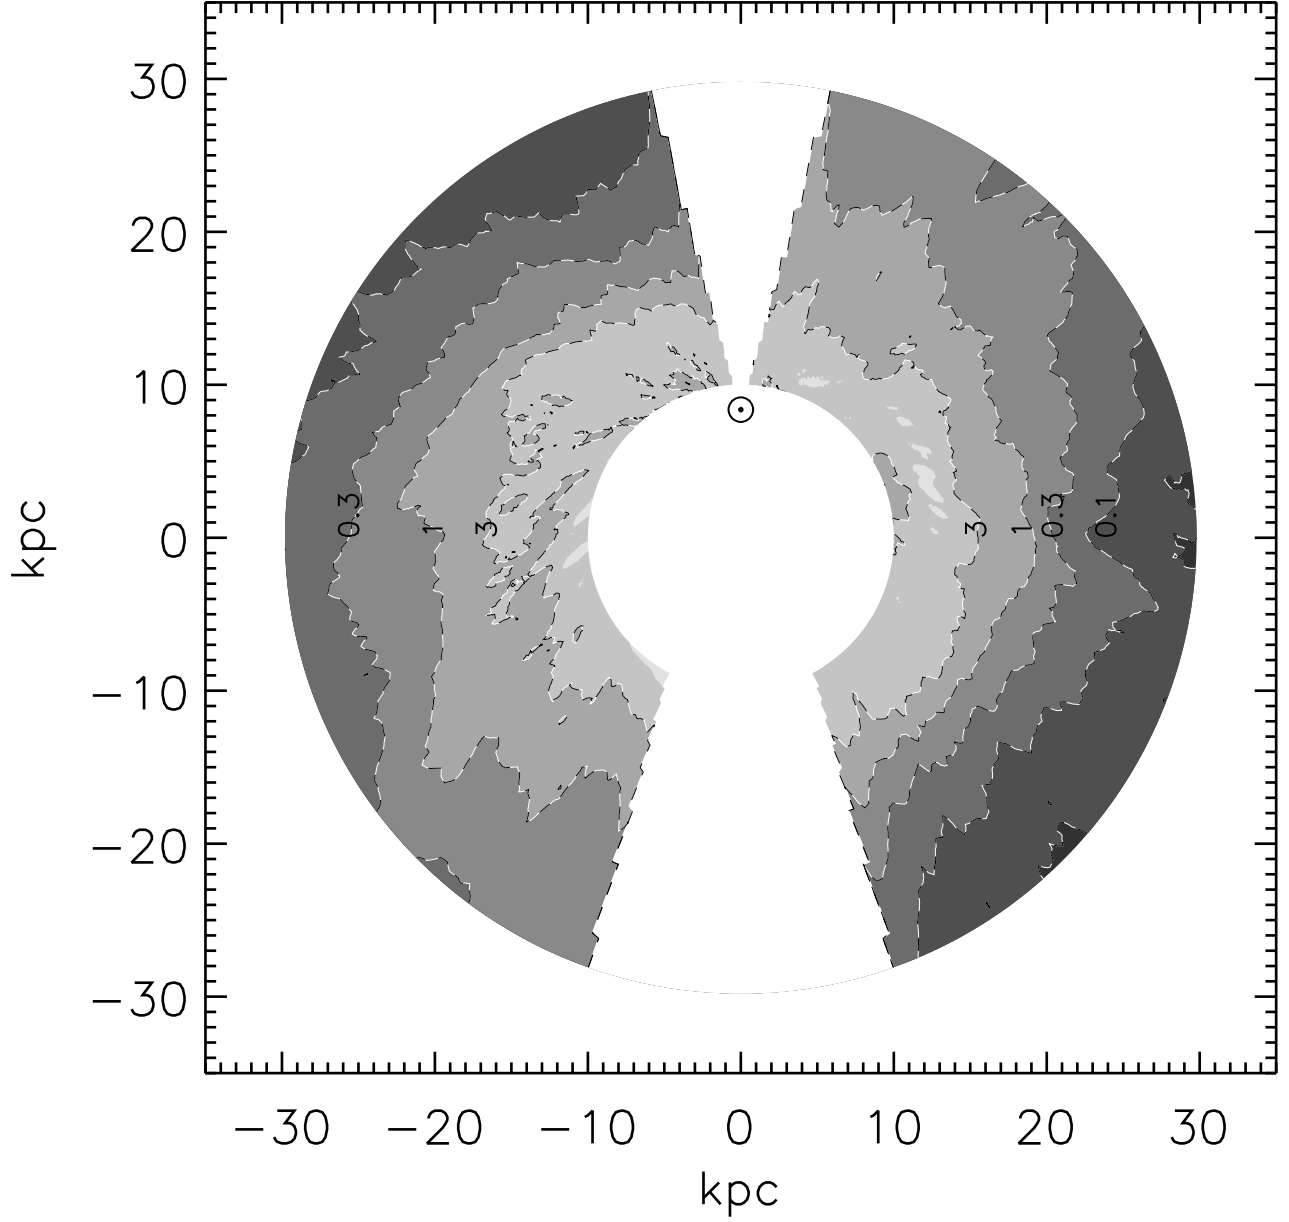
FIG. A17.— A contour map of $\Sigma(R, \phi)$ without any correction for elliptical orbits. Contours are evenly spaced in $\log_{10}$ space, at $(3, 1, 0.3, 0.1)$ M$_\odot$/pc^2.

From (B1), we construct the wavelet power spectrum, $|W_n(s)|^2$. We calculate the wavelet power spectrum for the dense set of frequencies, and then filter over a range of scales to find the power in a frequency band. The filtered power spectrum is given by

$$\overline{W_n}^2 = \frac{\Delta j \Delta \phi}{C_\delta} \sum_{j=j_1}^{j_2} \frac{|W_n(s_j)|^2}{s_j} \tag{B4}$$

where C_δ is a reconstruction factor that depends on the choice of wavelet. For the Morlet wavelet with $\omega_0 = 6$, $C_\delta = 0.776$.

REFERENCES

Arnal, E. M., Bajaja, E., Larrarte, J. J., Morras, R., & Pöppel,
    W. G. L. 2000, A&AS, 142, 35

Bajaja, E., Arnal, E. M., Larrarte, J. J., Morras, R., Pöppel,
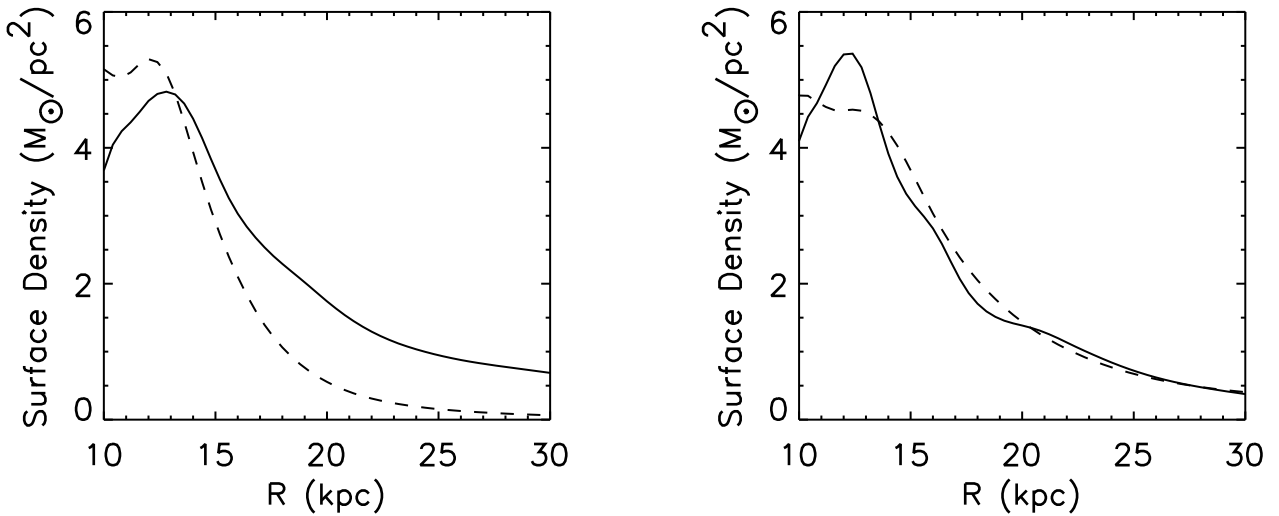   W. G. L., & Kalberla, P. M. W. 2005, A&A, 440, 767

Fig. A18.— Left: The average surface densities with no correction for $v_\Pi(R)$. The solid line is $\Sigma_{165}(R)$ and the dotted line is $\Sigma_{195}(R)$. Right: The average surface densities after being adjusted with $v_\Pi(R)$.

Battaner, E., Florido, E., & Sanchez-Saavedra, M. L. 1990, A&A, 236, 1

Binney, J. 1978, MNRAS, 183, 779

Binney, J., Jiang, I.-G., & Dutta, S. 1998, MNRAS, 297, 1237

Binney, J. & Merrifield, M. 1998, Galactic astronomy (Galactic astronomy / James Binney and Michael Merrifield. Princeton, NJ : Princeton University Press, 1998. (Princeton series in astrophysics) QB857 .B522 1998 ($35.00))

Blitz, L. & Spergel, D. N. 1991, ApJ, 370, 205

Bosma, A. 1991, in Warped Disks and Inclined Rings around Galaxies, 181–+

Brand, J. & Blitz, L. 1993, A&A, 275, 67

Briggs, F. H. 1990, ApJ, 352, 15

Burke, B. F. 1957, AJ, 62, 90

Burton, W. B. 1985, A&AS, 62, 365

—. The structure of our Galaxy derived from observations of neutral hydrogen (Galactic and Extragalactic Radio Astronomy), 295–358

Burton, W. B. & Te Lintel Hekkert, P. 1986, A&AS, 65, 427

Dehnen, W. 1998, AJ, 115, 2384

Djorgovski, S. & Sosin, C. 1989, ApJ, 341, L13

Edelsohn, D. J. & Elmegreen, B. G. 1997, MNRAS, 287, 947

Fleck, R. C. 1983, ApJ, 270, 507

Freudenreich, H. T., Berriman, G. B., Dwek, E., Hauser, M. G., Kelsall, T., Moseley, S. H., Silverberg, R. F., Sodroski, T. J., Toller, G. N., & Weiland, J. L. 1994, ApJ, 429, L69

García-Ruiz, I., Kuijken, K., & Dubinski, J. 2002a, MNRAS, 337, 459

García-Ruiz, I., Sancisi, R., & Kuijken, K. 2002b, A&A, 394, 769

Gum, C. S., Kerr, F. J., & Westerhout, G. 1960, MNRAS, 121, 132

Hartmann, D. & Burton, W. B. 1997, Atlas of Galactic Neutral Hydrogen (Atlas of Galactic Neutral Hydrogen, by Dap Hartmann and W. Butler Burton, pp. 243. ISBN 0521471117. Cambridge, UK: Cambridge University Press, February 1997.)

Henderson, A. P., Jackson, P. D., & Kerr, F. J. 1982, ApJ, 263, 116

Hunter, C. & Toomre, A. 1969, ApJ, 155, 747

Jiang, I. & Binney, J. 1999, MNRAS, 303, L7

Kahn, F. D. & Woltjer, L. 1959, ApJ, 130, 705

Kalberla, P. M. W., Burton, W. B., Hartmann, D., Arnal, E. M., Bajaja, E., Morras, R., & Pöppel, W. G. L. 2005, A&A, 440, 775

Kerr, F. 1968, in Stars and Stellar Systems, vol. 7, Nebulae and Interstellar Matter, ed. B. Middlehurst & L. Aller (Chicago: University of Chicago Press), 574–622

Kerr, F. & Westerhout, G. 1965, in Stars and Stellar Systems, Vol 5, Galactic Structure, ed. A. Blaauw & M. Schmidt (Chicago: University of Chicago Press), 167–202

Kerr, F. J. 1957, AJ, 62, 93

—. 1962, MNRAS, 123, 327

Kerr, F. J., Hindman, J. V., & Carpenter, M. S. 1957, Nature, 180, 677

Kuijken, K. & García-Ruiz, I. 2001, in ASP Conf. Ser. 230: Galaxy Disks and Disk Galaxies, ed. J. G. Funes & E. M. Corsini

Kuijken, K. & Tremaine, S. 1994, ApJ, 421, 178

Kulkarni, S. R., Heiles, C., & Blitz, L. 1982, ApJ, 259, L63

Levine, E. S., Blitz, L., & Heiles, C. 2006, to be submitted

López-Corredoira, M., Betancort-Rijo, J., & Beckman, J. E. 2002, A&A, 386, 169

Lozinskaya, T. A. & Kardashev, N. S. 1963, Soviet Astronomy, 7, 161

Marquardt, D. W. 1963, Journal of the Society for Industrial and Applied Mathematics, 11, 431

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in C. The art of scientific computing (Cambridge: University Press, —c1992, 2nd ed.)

Radhakrishnan, V. & Sarma, N. V. G. 1980, A&A, 85, 249

Sánchez-Salcedo, F. J. 2004, Journal of Korean Astronomical Society, 37, 205

Sparke, L. S. 1984, ApJ, 280, 117

—. 1995, ApJ, 439, 42

Sparke, L. S. & Casertano, S. 1988, MNRAS, 234, 873

Spicker, J. & Feitzinger, J. V. 1986, A&A, 163, 43

Toomre, A. 1983, in IAU Symp. 100: Internal Kinematics and Dynamics of Galaxies, 177–185

Torrence, C. & Compo, G. 1998, Bulletin of the American Meteorological Society, 79, 61

Voskes, T. 1999, M. Sc. Thesis, University of Leiden, astro-ph/0601653

Weinberg, M. D. 1998, MNRAS, 299, 499

Weinberg, M. D. & Blitz, L. 2005, ApJ, 0, 0

Westerhout, G. 1957, Bull. Astron. Inst. Netherlands, 13, 201

Wouterloot, J. G. A., Brand, J., Burton, W. B., & Kwee, K. K. 1990, A&A, 230, 21